Department of Computer & Information Science

Technical Reports (CIS)

University of Pennsylvania *Year* 2010

Security Weaknesses in the APCO Project 25 Two-Way Radio System

Sandy Clark* Perry Metzger† Zachary Wasserman‡
Kevin Xu** Matthew A. Blaze††

*University of Pennsylvania
†University of Pennsylvania
‡University of Pennsylvania
**University of Pennsylvania
††University of Pennsylvania, blaze@cis.upenn.edu

This paper is posted at ScholarlyCommons.
http://repository.upenn.edu/cis_reports/944

Security Weaknesses in the APCO Project 25 Two-Way Radio System

Sandy Clark, Perry Metzger, Zachary Wasserman, Kevin Xu, and Matt Blaze*
University of Pennsylvania

CIS Technical Report MS-CIS-10-34 – 18 November 2010

Abstract

APCO Project 25 ("P25") is a suite of wireless communications protocols designed for public safety two-way (voice) radio systems. The protocols include security options in which voice and data traffic can be cryptographically protected from eavesdropping. This report analyzes the security of P25 systems against passive and active attacks. We find a number of protocol, implementation, and user interface weaknesses that can leak information to a passive eavesdropper and that facilitate active attacks. In particular, P25 systems are highly susceptible to *active traffic analysis* attacks, in which radio user locations are surreptitiously determined, and *selective jamming* attacks, in which an attacker can jam specific kinds of traffic (such as encrypted messages or key management traffic). The P25 protocols make such attacks not only feasible but highly efficient, requiring, for example, significantly less aggregate energy output from a jammer than from the legitimate transmitters.

1 Introduction

APCO Project 25[15] (also called "P25") is a suite of digital protocols and standards designed for use in narrowband short-range (VHF and UHF) land-mobile wireless two-way communications systems. The system is intended primarily for use by public safety and other government users.

The P25 protocols are designed by an international consortium of vendors and users (centered in the United States), coordinated by the Association of Public Safety Communications Officers (APCO) and with its standards documents published by the Telecommunications Industry Association (TIA). Work on the protocols started in 1989, with new protocol features continuing to be refined and standardized on an ongoing basis.

The P25 protocols support both digital voice and low bit-rate data messaging, and are designed to operate in stand-alone short range "point-to-point" configurations or with the aid of infrastructure such as repeaters that can cover larger metropolitan and regional areas.

P25 supports a number of security features, including optional encryption of voice and data, based on either manual keying of mobile stations or "over the air" rekeying ("OTAR"[14]) through a key distribution center.

In this report, we examine the security of the P25 standards (and common implementations of it) against unauthorized eavesdropping, traffic analysis, active adjuncts to traffic analysis, and active interference.

*Author email: saender@cis.upenn.edu, pmetzger@cis.upenn.edu, zwass@seas.upenn.edu, kevinxu@seasupen.edu and blaze@cis.upenn.edu, respectively.

We find that these systems are strikingly vulnerable to a range of attacks. Practical attacks can leak information, including location information about members of a radio group, and can seriously mislead their users about the security state of their communication. The protocol is also highly susceptible to active denial-of-service, with *more than an order of magnitude less average power* required to effectively jam them than the analog systems they are intended to replace.

We describe an *active traffic analysis* attack that permits on-demand determination of the location of all of the users of a radio network, even when they are not actively using their radios. We also describe very low-energy *selective jamming attacks* that exploit a variety of protocol weakness, with the effect that encrypted users can be forced (knowingly or unknowingly) to revert to unencrypted mode. These attacks may also be used to entirely disable trunked mode communications for an entire radio network with strikingly low energy.

We begin with a brief overview of the P25 architecture, protocols, and usage model. We then describe fundamental weaknesses in the security protocols, implementations, and user interfaces that leak metadata, obscure security features from users, and that facilitate active attack. We then describe a technique to force all members of a radio group to periodically transmit in order to permit direction finding to locate their positions in space, with little chance of detection by the victims. Finally, we describe an active jamming architecture that exploits these weaknesses to prevent particular kinds of traffic, in a way that requires dramatically less average RF power to be emitted by the attacker than by the legitimate user. These attacks, which are difficult for the end-user to identify, can prevent encrypted traffic from being received and can force the users to disable encryption, or can be used to deny service altogether.

2 P25 Overview

P25 systems are intended as an evolutionary replacement for legacy analog FM narrowband two-way radio systems such as those used by local public safety agencies and national law enforcement and intelligence services. The systems are designed to be deployed without significant change to the user experience, radio channel assignments, spectrum bandwidth used, or network topology of conventional analog two-way radio. Users (or their vehicles) typically carry mobile transceivers[1] that receive voice communications from other users, with all radios in a group monitoring a common broadcast channel.

Mobile stations are equipped with "Push-To-Talk" buttons; the systems are half duplex, with at most one user transmitting on a given channel at a time. The radios typically either constantly receive on a single assigned channel or scan among multiple channels. Radios can be configured to mute received traffic not intended for them, and will ignore traffic for which a correct decryption key is not available.

P25 mobile terminal and infrastructure equipment is manufactured and marketed in the United States by a number of vendors, including E.F. Johnson, Harris, Icom, Motorola, RELM Wireless and Thales/Racal, among others. The P25 standards employ a number of patented technologies, including the voice codec, called IMBE[16]. Cross-licensing of patents and other technology is standard practice among the P25 equipment vendors, and so there are various features and implementation details common among equipment produced by different manufacturers. Motorola is perhaps the dominant U.S. vendor, and in this paper, we use Motorola's P25 product line to illustrate features, user interfaces, and attack scenarios.

For compatibility with existing analog FM based radio systems and for consistency with current radio spectrum allocation practices, P25 radios uses discrete narrow-band radio channels and not the spread spectrum techniques normally associated with digital wireless communication.

[1] Various radio models are designed be installed permanently in vehicles or carried as portable battery-powered "walkie-talkies".

Current P25 radio channels occupy a standard 12.5 KHz "slot" of bandwidth in the VHF or UHF land mobile radio spectrum. P25 systems use the same channel allocations as existing legacy narrowband analog FM two-way radios. To facilitate a gradual transition to the system, the standard requires that P25-compliant radios be capable of operating in the legacy analog mode as well as digital. Legacy analog radios cannot, of course, demodulate digital P25 transmissions, which are received as only as buzzing static on conventional FM receivers, but current P25 radios can be configured to demodulate analog transmissions and transmit in the legacy analog mode.

In the currently deployed C4FM modulation scheme, the 12.5kHz channel is used to transmit a four-level signal, sending two bits with each symbol at a rate of 4800 symbols per second, for a total bit rate of 9600bps.[2]

P25 radio systems can be configured in three different network topologies, requiring varying degrees of infrastructural support in the area of coverage:

- In the basic configuration, called *simplex,* the members of a group all set their transmitters and receivers to a single common frequency with each mobile station receiving transmissions from others and broadcasting its own transmissions on that frequency. The range of a simplex system is the area over which each station's transmissions can be received directly by the other stations, which is limited by terrain, power level, and interference from co-channel users.

- In another configuration, *repeater operation,* mobile stations transmit on one frequency to a fixed-location repeater, which in turn retransmits communications on a second frequency on which all the mobiles in a group receive. Repeater configurations thus use two frequencies per channel. The repeater typically possesses both an advantageous geographical location and access to electrical power. Repeaters extend the effective range of a system by rebroadcasting mobile transmissions at higher power and from a greater height.

- In a third configuration, called *trunking,* mobile stations transmit and receive on a variety of frequencies as orchestrated by a "control channel" supported by a network of base stations. By dynamically allocating transmit and receive frequencies from among a set of allocated channels, scarce radio bandwidth may be effectively time and frequency domain multiplexed among multiple groups of users.

For simplicity, this report focuses chiefly on weaknesses and attacks that apply to all three configurations, although we will briefly discuss a denial of service attack specific to P25 trunked configurations as well.

As P25 is a digital protocol, it is technically straightforward to encrypt voice and data traffic, something that is far more difficult in the analog domain systems it is designed to replace. However, encryption is an optional feature, and even radios equipped for encryption have the capability to operate in the clear mode. Keys may be manually loaded into mobile units or may be updated at intervals using the Over The Air Rekeying ("OTAR") protocol.

P25 also provides for a low-bandwidth data stream that piggybacks atop voice communications, and for a higher bandwidth data transmission mode in which data is sent independent of voice. (It is this facility which enables the OTAR protocol, as well as attacks we describe below to actively locate mobile users.)

[2] As has been mentioned, this so-called "Phase 1" scheme is designed to co-exist with analog legacy systems. A quadrature phase shift keying system has also been specified that permits similar bandwidth using only 6.25kHz of spectrum, and a future "phase 2" system will use either TDMA or FDMA and a more efficient vocoder to again double channel capacity by permitting two simultaneous users in each 6.25kHz slot. However, Phase 2 systems have not yet been widely deployed, and in any case do not affect the security analysis in this report.

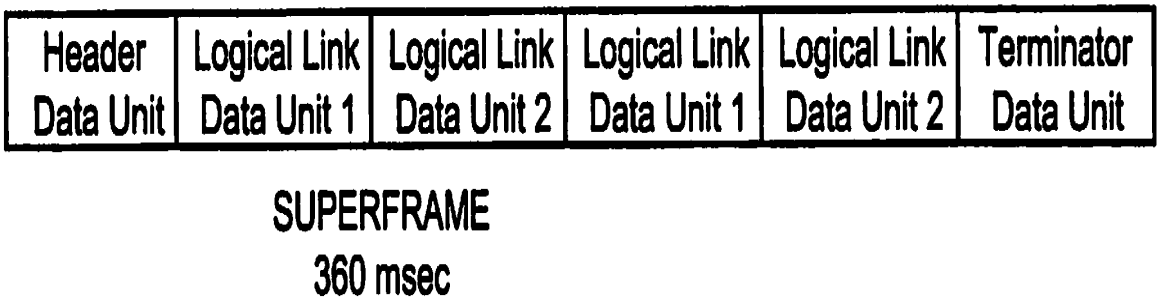

Figure 1: Data Frame structure (from Project 25 FDMA - Common Air Interface: TIA-102.BAAA-A)

2.1 The P25 Protocols

This section is a brief overview of the most salient features of the P25 protocols relevant to rest of this paper. The P25 protocols are quite complex, and the reader is urged to consult the standards themselves for a complete description of the various data formats, options, and message flows. An excellent overview of the most important P25 protocol features can be found in reference [4].

The P25 Phase 1 (the currently deployed version) RF-layer protocol uses a four level code over a 12.5kHz channel, sending two bits per transmitted symbol at 4800 symbols per second or 9600 bits per second.

A typical transmission consists of a series of *frames,* transmitted back-to-back in sequence. The start of each frame is identified by a special 24 symbol (48 bit) frame synchronization pattern.

This is immediately followed by a 64 bit field containing 16 bits of information and 48 bits of error correction. 12 of these bits identify the network on which the message is being sent (this is the "NAC" field) – a radio remains muted unless a received transmission contains the correct NAC, which prevents unintended interference by distinct networks using the same set of frequencies. 4 of the bits identify the type of the frame – these are the so-called "Data Unit ID" or "DUID". These bits identify a frame as a voice transmission header, as one of two kinds of voice frame, a voice transmission trailer, a data packet, or a frame associated with trunked system operation. Most frame types are of fixed length, the exception being packet data frames.

In a voice transmission, a header will be followed by a sequence of voice "superframes" (called "LDUs" and discussed below), and finally a trailer ("terminator") frame. See Figure 1

Headers frames contain a 16 bit field designating the destination talk group ("TGID") for which a transmission is intended. This permits radios to mute transmissions not intended for them. The header also contains information for use in encrypted communications, specifically an initialization vector (designated the "Message Indicator" or "MI" in P25, which is 72 bits wide but effectively only 64 bits), an eight bit Algorithm ID, and a 16 bit Key ID. Transmissions in the clear set these fields to all zeros. This information is also accompanied by a large number of error correction bits.

The actual audio payload, encoded as IMBE voice subframes, is sent inside *Link Data Units* (LDUs). A voice LDU contains a header followed by a sequence of nine 144 bit IMBE voice subframes (each of which encodes 20ms of audio, for a total 180ms of encoded audio in each LDU frame), plus additional metadata and a small amount of piggybacked low speed data. Each LDU, including headers, metadata, voice subframes, and error correction is 864 symbols (1728 bits) long.

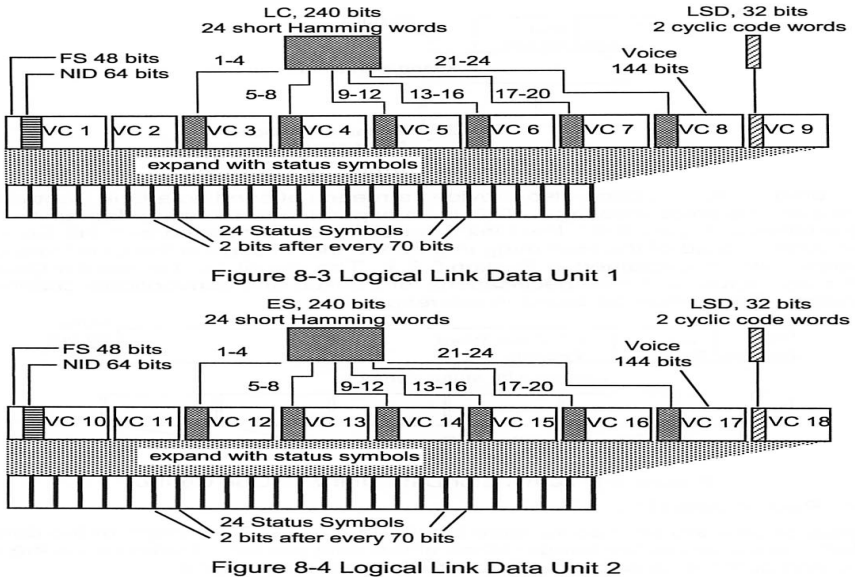

Figure 2: Logical Data Unit structure (from Project 25 FDMA - Common Air Interface: TIA-102.BAAA-A)

A voice transmission thus consists of a header frame followed by an arbitrary length alternating sequence of LDU frames in two slightly different formats (called LDU1 and LDU2 frames, which differ in the metadata they carry), followed by a terminator frame. Note that the number of voice LDU1 and LDU2 frames to be sent in a transmission is not generally known at the start of the transmission, since it depends on how long the user speaks.

LDU1 frames contain the source unit ID of a given radio (a 24 bit field), and either a 24 bit destination unit ID (for point to point transmissions) or a 16 bit TGID (for group transmissions).

LDU2 frames contain new MI, Algorithm ID and Key ID fields.

Because all the metadata required to recognize a transmission is available over the course of two LDU frames, a transmission may be "caught up with" from the middle without the need to have received the header frame. Voice LDU frames alternate between the LDU1 and LDU2 format. An LDU1/LDU2 pair is also called a "superframe", since once both are received a receiver can begin processing the transmission even if the initial transmission header was missed.

See Figure 2 for the structure of the LDU1 and LDU2 frames.

Terminator units, which may follow either an LDU1 or LDU2 frame, indicate the end of a transmission.

A separate format exists for (non-voice) packet data frames. Data frames may optionally request acknowledgment to permit immediate retransmission in case of corruption. A header, which is always unencrypted, indicates which unit ID has originated the packet or is its target. (These features will prove important in the discussion of active radio localization attacks.)

Trunking systems also use a frame type of their own on their control channel. (We do not discuss the details of this frame type, as they are not relevant to our study.)

It is important to note a detail of the error correction codes used for the voice data in LDU1 and LDU2 frames. The IMBE codec has the feature that not all bits in the encoded representation are of equal importance in regenerating the original transmitted speech. To reduce the amount of error correction needed in the frame, bits that contribute more to intelligibility receive more error correction than those that contribute less, with the least important bits receiving no error correction at all. Although this means that the encoding

of voice over the air is more efficient, it also eliminates the possibility of protecting voice transmissions with block ciphers or message authentication codes, as we explain below.

2.2 Security Features

P25 provides options for traffic confidentiality using symmetric-key ciphers, which can be implemented in software or hardware. The standard supports mass-market "Type 2/3/4' crypto engines (such as DES and AES) for unclassified domestic and export users, as well as NSA-approved "Type 1" cryptography for government classified traffic. (The use of Type 1 hardware is tightly controlled and restricted to classified traffic only; even sensitive criminal law enforcement surveillance operations typically must use commercial Type 2/3/4 cryptography.)

The DES, 3DES and AES ciphers are specified in the standard, in addition to the null cipher for cleartext. The standard also provides for the use of vendor-specific proprietary algorithms (such as 40 bit RC4 for radios aimed at the export market). [10]

At least for Type 2, 3 and 4 cryptography, pre-shared symmetric keys are used for all traffic encryption. The system requires a key table located in each radio mapping unique Key ID+Algorithm ID tuples to particular symmetric cipher keys stored within the unit. This table may be keyed manually or with the use of an Over The Air Rekeying protocol. A group of radios can communicate in encrypted mode only if all radios share a common key (labeled with the same Key ID).

Many message frame types contain a tuple consisting of an initialization vector (called a "Message Indicator" or "MI" in the protocol standard), a Key ID and an Algorithm ID. A clear transmission is indicated by a zero MI and KID and a special ALGID. The key used by a given radio group may thus change from message to message and even from frame to frame (some frames may be sent encrypted while others are sent in the clear).

Because of the above-described property of the error correction mechanisms used, especially in voice frames such as the LDU1 and LDU2 frame types, there is no mechanism to detect errors in certain portions of transmitted frames. This is a deliberate design choice, with the voice encoding designed to permit undetected corruption of portions of the frame that are less important for intelligibility.

This error-tolerant design means that standard block cipher modes (such as Cipher Block Chaining) cannot be used for voice encryption; block ciphers require the accurate reception of an entire block in order for any portion of the block to be correctly decrypted. P25 voice encryption therefore uses stream ciphers, in which a cryptographic keystream generator produces a pseudorandom bit sequence that is XORd with the data stream to encrypt (on the transmit side) and decrypt (on the receive side). In order to permit conventional block ciphers (including DES and AES) to be used as stream ciphers, they are run in Output Feedback mode ("OFB")) in order to generate a key stream. (Some native stream ciphers, such as RC4, have been implemented by some manufacturers as well, particularly for use in export radios limited to short key lengths.)

For the same reason – received frames must tolerate the presence of some bit errors – cryptographic message authentication codes ("MACs"), which fail if any bit errors whatsoever are present, are not used.

As noted above, cryptographic keying of P25 radios may be done manually by a technician with a special keyloading device (called a *Key Variable Loader* or KVL) or may be performed remotely via the OTAR protocol. The OTAR protocol relies on each mobile having pre-shared unit-specific keying material (key encrypting keys) that permit a remote Key Management Facility (or "KMF") to securely add, update, and remove elements of the radios' traffic key tables.

3 Security Analysis

In the previous section, we described a highly ad hoc, constrained architecture that, we note, departs in significant ways from conservative security design, does not provide clean separation of layers, and lacks a clearly stated set of requirements against which it can be tested.

This is true even in portions of the archiecture, such as the packet data frame subsystem, which are at least in theory compatible with well understood standard cryptographic protocols, such as those based on block ciphers and MACs.

This ad hoc design by itself represents a security concern, and could be considered a basic architectural weakness. In fact, the design introduces significant certificational weaknesses in the cryptographic protection provided.

These weaknesses do not, in and of themselves, automatically result in exploitable vulnerabilities. However, they weaken and complicate the guarantees that can be made to higher layers of the system. Given the overall complexity of the P25 protocol suite, and especially given the reliance of upper layers such as the OTAR subsystem on the behavior of lower layers, such deficiencies make the security of the overall system much harder for a defender to analyze.

The P25 implementation and user interfaces, too, suffer from an ad hoc design that, we shall see, does not fare well against an adversarial threat. There is no evidence in the standards documents, product literature, or other documentation of user interface or usability requirements, or of testing procedures such as "red team" exercises or user behavior studies.

As we shall see later in this paper, taken in combination, the design weaknesses of the P25 security architecture and the standard implementations of it admit practical, exploitable vulnerabilities that routinely leak sensitive traffic and that allow an active attacker remarkable leverage.

3.1 Protocol Weaknesses

At the root of many of the most important practical vulnerabilities in P25 systems are a number of fundamentally weak cryptographic, security protocol, and coding design choices.

3.1.1 Unauthenticated Message Traffic

Because no MACs are employed on voice and most other traffic, even in encrypted mode, it is trivial for an adversary to masquerade as a legitimate user, to inject false voice traffic, and to replay captured traffic, even when all radios in a system have encryption configured and enabled.

The ability for an adversary to inject false traffic without detection is, of course, a fundamental weakness by itself, but also something that can be serve as a stepping stone to more sophisticated attacks.

3.1.2 Unprotected Metadata

Even when encryption is used, much of the basic metadata that identifies the systems, talk groups, sender and receiver user IDs, and message types of transmissions are sent in the clear and are directly available to a passive eavesdropper for traffic analysis and to facilitate other attacks. While some of these fields can be optionally encrypted (although the use of encryption is not tied to whether voice encryption is enabled), others must always be sent in the clear due to the basic architecture of P25 networks.

For example, the start of every frame of every transmission includes a *Network Identifier* ("NID") field that includes the 12 bit Network Access Code (NAC) and the 4 bit frame type ("Data Unit ID"). The NAC

code identfies the network on which the transmission is being sent; on frequencies that carry traffic from multiple networks, it effectively identifies the organization or agency from which a transmission originated. The Data Unit ID identifies the type of traffic, voice, packet data, etc. Several aspects of the P25 architecture requires that the NID be sent in the clear. For example, repeaters and other infrastructure (which do not have access to keying material) use it to control the processing of the traffic they receive. The effect is that the NAC and type of transmission is always available to a passive adversary on every transmission.

For voice traffic, a *Link Control Word* ("LCW") is included in every other LDU voice frame (specifically, in the LDU1 frames). The LCW includes the transmitter's unique unit ID (somewhat confusingly called the "Link IDs" in various places in the standard). The ID fields in the LCW can be optionally encrypted, but whether they are actually encrypted is not intrinsically tied to whether encryption is enabled for the voice content itself (rather it is indicated by a "protected" bit flag in the LCW).

A widely deployed implementation error exacerbates the unit ID information leaked in the LCW. We found that in the Motorola's (and possibly other vendors') P25 products we examined, the LCW protection bit appears never to be set; the option to encrypt the LCW is never enabled, even when the voice traffic itself was otherwise encrypted. That is, in every P25 product of which we are aware, the sender's unit Link ID is always sent in the clear, even for encrypted traffic. This, of course, greatly facilitates traffic analysis of encrypted networks by a passive adversary; it also simplifies the certain active attacks we discuss in a subsequent section.

3.1.3 Use of stream ciphers

A well known weakness of stream ciphers is that attackers who know the plaintext content of any encrypted portion of transmission may make arbitrary changes to that content at will simply by flipping appropriate bits in the data stream. For this reason, it is generally recommended that stream ciphers always be used in conjunction with MACs, but the same design decision (error tolerance) that forced the use of stream ciphers in P25 also prevents the use of MACs.

3.1.4 ECC weaknesses

Error correction codes in P25 frame formats are highly optimized for the various kinds of content in the frame. That is, a single error correcting code is not used across the entire frame. Instead, different sections of P25 frames are error corrected in independent ways, with separate codes providing error correction for relatively small individual portions of the data stream. This leaves the frames vulnerable to highly efficient active jamming attacks that target only small portions of frames – see the section on "Denial of Service" for details.

3.2 Implementation and Usability Weaknesses

P25 mobile radios are intended to support a range of government and public safety applications, many of which, such as covert law enforcement surveillance, require both a high degree of confidentiality as well as usability and reliability.

As noted above, the security features of P25 radios assume a centrally-controlled key distribution infrastructure shared by all users in a system. Once cryptographic keys have been installed in the mobile radios, either by a manual key loading device or through OTAR, the radios are intended to be simple to operate in encrypted mode with little or no interaction from the user.

Unfortunately, the security features are often difficult to use reliably in practice. In this section, we focus on examples drawn from Motorola's P25 product line. Motorola is a major vendor of P25 equipment in the United States and elsewhere, supplying P25 radios to the federal government as well as state and local agencies. Other vendors' radios have similar features; we use the Motorola products strictly for illustration. In particular, we set up a small encrypted P25 network using a set of Motorola Model XTS-5000 handheld radios.

3.2.1 User Interface Weaknesses

Most P25 radios are highly configurable. Indeed, most vendors do not impose a single standard user interface, but rather allow the radio's buttons, switches and "soft" menus to be customized by the customer. While this may seem an advantageous feature that allows each customer to configure its radios to best serve its application, the effect of this highly flexible design is that any given radio's user interface is virtually guaranteed to have poorly documented menus, submenus and button functions.

Because the radios are customized for each customer, the manuals are often confusing and incomplete when used side-by-side with an end-user's actual radio. For example, the Motorola XTS-5000 handheld P25 radio's manual[11] consists of nearly 150 pages that describe dozens of possible configurations and optional features, with incomplete instructions on how to activate features and interpret displayed information that typically advise the user to check with their local radio technician to find out how a given feature or switch works. (Other manufacturers' radios have a similarly configurable design). That is, every customer must, in effect, produce a custom user manual that describes how to properly use the security features as they happen to have been configured.

In a typical configuration for the XTS5000, outbound encryption is controlled by a rotating switch located on the same stem as the channel selector knob. We found it to be easy to accidentally turn off encryption when switching channels. And other than a small picture etched on this switch, there is little positive indication that the radio is or is not operating in encrypted mode.

On the XTS portable radios, a flashing LED indicates the reception of encrypted traffic. However, the same LED serves multiple purposes. It glows steady to indicate transmit mode, "slow" flashes to indicate received cleartext traffic, a busy channel, or low battery, and "fast" flashes to indicate received encrypted traffic. We found it to be very difficult to distinguish reliably between received encrypted traffic and received unencrypted traffic. Also, the LED and the "secure" display icon are likely out of the operator's field of view when an earphone or speaker/microphone is used or if the radio is held up to the user's ear while listening (or mouth when talking).

The Motorola radios can be configured to give an audible warning of clear transmit or receive in the form of a "beep" tone sounded at the beginning of each outgoing or incoming transmission. But the same tone is used to indicate other radio events, including button presses, low battery, etc, and the tone is difficult to hear in noisy environments.

In summary, we found it to be quite easy on the XTS-5000 to accidentally transmit in the clear, and correspondingly difficult to determine whether an incoming message was encrypted or with what key.

3.2.2 Clear Traffic Accepted in Encrypted Mode

All models of P25 radios of which we are aware will receive any traffic sent in the clear even when they are in encrypted mode. There is no configuration option to reject or mute clear traffic. While this may have some benefit to ensure interoperability in emergencies, it also means that a user who mistakenly places the "secure" switch in the "clear" position is unlikely to detect the error.

Because it is difficult to determine that one is receiving an accidentally non-encrypted signal, messages from a user unintentionally transmitting in the clear will still be received by all group members (and anyone else eavesdropping on the frequency), who will have no indication that there is a problem unless they happen to be actively monitoring their receivers' displays during the transmission.

Especially in light of the user interface issues discussed above, it is remarkably easy for some or all of an encrypted radio group to mistakenly operate in the clear without this being noticed for some time. If a subset of encrypted users are accidentally set for clear mode, the other users will still hear them. And as long as the clear users have the correct keys, they will still hear the encrypted transmissions even while their radios are inadvertently set to transmit in the clear.

3.2.3 Cumbersome Keying

P25 key management is designed for centralized control. As noted above, in most secure P25 products (including Motorola's), key material is loaded into radios either via a special key loader that is physically attached to the radio or through the OTAR protocol (to a KMF server that reached through the radio network).

There is no provision for individual groups of users to create ad hoc keys for short term or emergency use when they find that some members of a group lack the key material held by the others. That is, there is no mechanism for peers to engage in public key negotiation among themselves over the air or for keys to be entered into radios by hand without the aid of a keyloader.

This means that in there is no way for most users to themselves add a new member to the group or to recover if a radio is discovered to be missing the key during a sensitive operation. In systems that use automatic over-the-air keying at regular intervals, this can be especially problematic. If common keys get "out of sync" after some users have updated keys before others have, all users must revert to clear mode for the group to be able to communicate.[3]

3.3 Discussion

As we have seen, the P25 protocols and its implementations suffer from a range of weaknesses that, taken individually, might represent only relatively small risks or that can be effectively mitigated with careful radio configuration and user vigilance. But taken togther, they interact in far more powerful ways.

For example, if users are accustomed to occasionally having keys be out of sync and must frequently switch to clear mode, the risk that a user's radio will mistakenly remain in clear mode even when keys are available increases greatly.

More seriously, these vulnerabilities provide a large menu of options that increase the leverage for targeted active attacks that become far harder to defend against. In the following sections, we discuss several attacks against P25 systems that exploit combinations of these protocol, implementation and usability weaknesses to extract sensitive information, deny service, or manipulate user behavior in encrypted P25 systems.

4 Traffic Analysis

"Traffic Analysis" is the technique of inferring information from intercepted transmissions even when their contents are encrypted or largely encrypted and thus opaque to the eavesdropper. The P25 protocol pro-

[3] This scenario is a sharp counterexample to the cryptographic folk wisdom that frequently changing keys yields more security.

vides a number of traffic analysis opportunities to an interceptor, several of which are not possible or are substantially more difficult in ordinary analog systems.

Many kinds of traffic analysis are possible and it is beyond the scope of this document to discuss them all. Instead, we focus on a single scenario, that of an adversary who is trying to learn about or evade a (P25 radio-equipped) surveillance team that may be following her. Traffic analysis can be used to confirm this suspicion, even when the surveillance traffic is encrypted.

4.1 Passive Location Tracking

Transmitting radio sources are is generally susceptible to geolocation through direction finding and triangulation techniques. With the proper equipment, direction finding is possible for virtually any radio signal format, whether it is analog or digital. However, the P25 protocol provides particularly valuable addressing information to an attacker that is not typically available in legacy analog systems.

An eavesdropper familiar with the frequencies used by a given agency may readily listen to that frequency set and determine which group IDs are regularly in use, and may employ direction finding equipment to locate the radios corresponding to a particular group. Group IDs are always sent in the clear.

For most traffic a passive eavesdropper can track individual radios simply by noting the senders'Unit Link ID numbers sent in the clear in various metadata fields during transmissions[4].

In encrypted mode, Unit Link ID numbers can be optionally protected in voice frames (but not for data frames). However, as noted in the previous section, this feature is apparently not implemented in many vendors' equipment, even when encryption is used to protect content. And for packet data messages (for example, when the OTAR protocol is used for key management), the protocol specifies that Unit Link IDs are always transmitted in the clear in the data frame's header block even when the packet data itself is encrypted.

If only voice frames are sent on a given encrypted link, Unit Link ID numbers may not be visible in the clear if the radios correctly implement the "protected" flag in the LCW (which, we note, is often not implemented). However, even without knowing the link IDs, valuable information may be easily obtained. For example, an entirely passive eavesdropper could use direction finding to discover whether the movements of members of a surveillance team correlate wellwith his own movements in public places.

An active adversary has even richer traffic analysis options.

4.2 Active Location Tracking

Generally, a radio's location may be tracked only if it is actively transmitting. P25 provides a convenient means for an attacker to induce otherwise silent radios to transmit, permitting active continuous tracking of a radio's user. [8, 7]

The P25 protocol provides for a packet data transmission system. Packets may be sent in either an unconfirmed mode, in which retransmission in the event of errors is handled by a higher layer of the protocol, or in confirmed mode, in which the destination radio must acknowledge successful transmission of a data frame or request that it be retransmitted.

If the Unit Link IDs of a radio group are known to an adversary, she may periodically send out intentionally corrupted data frames to each member of the group. Only the header CRCs need check cleanly for a data frame to be replied to – all content CRCs may be intentionally defective. The group member radios

[4] As discussed in previous sections, each P25 radio in a given system is assigned a unique 24 bit unit ID number, called variously the "Unit ID" or the "Link ID", that is transmitted along with all outbound traffic. To avoid ambiguity, we will refer to this unique identifier as the *Unit Link ID* in this section.

will then reply requesting retransmission of this defective data frame, and may be tracked based on this transmission. It is unlikely that such corrupted data frames will be noticed, especially since the frames will always be rejected as corrupt before being passed to the higher layers in the radio's software.

While we are unaware of any systems that will refuse to respond to a data frame that is not properly encrypted, even if encryption is enabled and a radio refuses to pass unencrypted frames to higher level firmware, the attacker may easily construct a forged but valid encryption auxiliary header simply by capturing legitimate traffic and inserting a stolen encryption header. This is possible because on receiving a corrupt packet, there is no expectation that the packet will decrypt correctly, and therefore, the victim radio will respond requesting retransmission, acting as an oracle for its presence and also allowing direction finding on demand.

If the target radios' Unit Link IDs are for some reason unknown to the attacker, she may straightforwardly attempt a "wardialing" [12] attack in which she systematically guesses Unit Link IDs and sends out requests for replies, taking note of which ID numbers respond. However, in a trunked system or a system using Over the Air Rekeying, or in a system where members of the radio group occasionally transmit voice in the clear, Link IDs will be readily available without resorting to wardialing in this manner.

With this technique, an attacker can easily "turn the tables" on covert users of P25 mobile devices, effectively converting their radios into location tracking beacons.

5 Denial of Service

P25 uses a narrowband (C4FM in Phase 1 systems) modulation scheme designed to fit into channels compatible with the current spectrum management practices for two-way land mobile radio. Unfortunately, although this was a basic design constraint, it not only denies P25 systems the jamming resistance of digital spread spectrum systems, it actually makes them *more* vulnerable to denial of service than the analog systems they replace. The P25 protocols also permit potent new forms of deliberate interference, such as *selective attacks* that induce security downgrades, a threat that is exacerbated by usability deficiencies in current P25 radios.

5.1 Jamming in Analog and Digital Systems

Jamming attacks, in which a receiver is prevented from successfully interpreting a signal by noise injected onto the over the air channel, are a long-known and widely studied problem in wireless systems.

In ordinary narrowband channelized analog FM systems, jamming and defending against jamming is a matter of straightforward analysis. The jammer succeeds when it overcomes the power level of the legitimate transmitter at the receiver. Otherwise the "capture effect", a phenomenon whereby the strongest of two signals at or near the same frequency is the one demodulated by the receiver, permits the receiver to continue to understand the transmitted voice signal. An attacker may attempt to inject an intelligible signal or actual noise to prevent reception. In practice, an FM narrowband jammer will succeed reliably if it can deliver 3 to 6 db more power to the receiver than the legitimate transmitter (to exceed the "capture ratio" of the system). Jamming in narrowband systems is thus for practical purposes a roughly equally balanced "arms race" between attacker and defender. Whoever has the most power wins.[5]

[5] As a practical matter, the analog jamming arms race is actually tipped slightly in favor of the *defender,* since the attacker generally also has to worry about being discovered (and then eliminated) with radio direction finding and other countermeasures. More power makes the jammer more effective, but also easier to locate.

In digital wireless systems, the jamming arms race is more complex, depending on the selected modulation scheme and protocol.

Spread spectrum systems [3], and especially direct sequence spread spectrum systems, can be made robust against jamming, either by the use of a secret spreading code or by more clever techniques described in [6, 1]. Without special information, a jamming transmitter must increase the noise floor not just on a single frequency channel, but rather across the entire band in use, at sufficient power to prevent reception. This requires far more power than the transmitter with which it seeks to interfere, and typically more aggregate power than an ordinary transmitter would be capable of. Modern spread spectrum systems can enjoy an average power advantage of 30db or more over a jammer. That is, in a spread spectrum system operating over a sufficiently wide band, a jammer can be forced to deliver more than 30db more aggregate power to the receiving station than the legitimate transmitter.

By contrast, in a narrow-band digital modulation scheme such as P25's current C4FM mode (or the lower-bandwidth Phase 2 successors proposed for P25), jamming requires only the transmission of a signal at a level near that of the legitimate transmitter. Competing signals arriving at the receiver will prevent clean decoding of a transmitted symbol, effectively randomizing or setting the received symbol. [2] That is, C4FM modulation suffers from approximately the same inherent degree of susceptibility to jamming as narrowband FM – a jammer must simply deliver slightly more power to the receiver than the legitimate transmitter.

But, as we will see below, the situation is actually far more favorable to the jammer than analysis of its modulation scheme alone might suggest. In fact, the *aggregate* power level required to jam P25 traffic is actually much *lower* than that required to jam analog FM. This is because an adversary can disrupt P25 traffic very efficiently by targeting only specific small portions of frames to jam and turning off its transmitter at other times.

5.2 Partial Frame Jamming Attacks

We found that the P25 protocols are vulnerable to highly efficient jamming attacks that exploit not only the narrowband modulation scheme, but also the structure of the transmitted messages.

Most P25 frames contain one or more small metadata subfields that are critical to the interpretation of the rest of the frame. For example, if the 4-bit Data Unit ID, present at the start of every frame, is not received correctly, receivers cannot determine whether it is a header, voice, packet or other frame type. This is not the only critical subfield in a frame, but it is illustrative for our purposes.

It is therefore unnecessary for an adversary to jam the entire transmitted data stream in order to prevent a receiver from receiving it. It is sufficient for an attacker to prevent the reception merely of those portions of a frame that are needed for the receiver to make sense of the rest of the frame.

Unfortunately, the P25 frame encoding makes it particularly easy and efficient for a jammer to attack these subfields in isolation.

For example, a voice frame is 1728 bits in length. The entire *NID* subfield containing the NAC + DUID (and its error correction code) represents only 64 bits of these 1728 bits. Jamming just the 64 bit NID subfield effectively denies the receiver the ability to interpret the other 1664 bits of the frame, even if those bits are received unmolested . A jammer synchronized to attack just the NID subfield of voice transmission would need to operate at a duty cycle of only 3.7% during transmissions. Such a pulse lasts only about 1/100th of a second.

To efficiently jam particular frame subfields, a jammer must synchronize its transmissions so that it begins transmitting at or just before the the first symbol of the targeted field is sent by the transmitter under attack, and end just after the last symbol of the field has been sent. At 4800 symbols per second, each symbol lasts just longer than 0.2ms. This may seem at first to require an impossibly high degree of timing

synchronization. But the P25 framing scheme actually makes it quite straightforward for a jammer equipped with its own receiver to tightkly synchronize to the target transmitter. Recall that each frame begins with an easily-recognized frame synchronization word, which the jammer can use to precisely trigger its interference so that it begins and ends at exactly the desired symbols.

By careful synchronization, a jammer that attacks only the NID subfield of voice traffic can reduce its overall energy output so that it effectively has *more than 14db of average power advantage* over the legitimate transmitter.

It may be possible to improve the advantage to the jammer even more by careful analysis of the error correction codes used in particular subfields in order to reduce the number of bits in the subfield that have to be jammed. (We assumed conservatively above that the attacker must jam *every* bit of the 64 bit NID field in order to prevent correct reconstruction of at least one bit of the NID payload, which clearly can be improved upon). This would permit even lower transmission times and average emitted power. It is only necessary to reliably (though not necessarily perfectly) prevent correct interpretation of a single critical protocol field by the receiver – it is not necessary to fully obliterate it.

Properly synchronized, a P25 jamming system can operate at a very low duty cycle that not only saves energy at the jammer and makes its equipment smaller and less expensive, but also makes the existence of the attack difficult to diagnose and detect, and, if detected, require the use of specialized equipment to locate it. Such a jamming system need only be relatively inexpensive, requires only a modest power supply, and is trivial to deploy in a portable configuration that carries little risk to the attacker, as described below.

We note that there is no analogous low-duty cycle jamming attack possible against the narrowband FM voice systems that P25 replaces.

5.3 Selective Jamming Attacks

An attacker need not attempt to jam every transmitted frame. The attacker can pick and choose which frames to attack in order to encourage the legitimate users to alter their behavior in particular ways.

For example, it is straightforward to monitor for a non-zero MI field in a header frame (indicating an encrypted transmission) and to selectively jam portions of subsequent frames, while leaving clear transmissions alone, in order to create the impression to the users of a radio network that, for unknown technical reasons, encryption has malfunctioned while clear transmission remains viable, thus inducing the users to downgrade to clear transmissions. If the users are already conditioned (through other weaknesses in P25) to unreliable cryptography, such an attack might be dismissed as routine.

As another possibility, an attacker could choose to attack only uplink messages on the control channel of a trunked P25 system, thus effectively denying use of the entire trunked network at an extremely low cost to the attacker.

In addition to the complexities of detecting and direction-finding an attack lasting mere hundredths or even thousandths of a second, adversaries can take steps to render their attacks less vulnerable to detection and more difficult for the operators of a radio network to prevent. For example, an attacker could choose to deploy multiple battery operated jamming devices in a metropolitan area, placing them in public locations to make tracing of the devices harder, or even surreptitiously attaching them to the vehicles of third parties such as taxis or delivery trucks to cause confusion, and to make the jammers harder to locate. Such devices may be made arbitrarily programmable, changing which of a group of devices is active at any one time or even taking commands over the air.

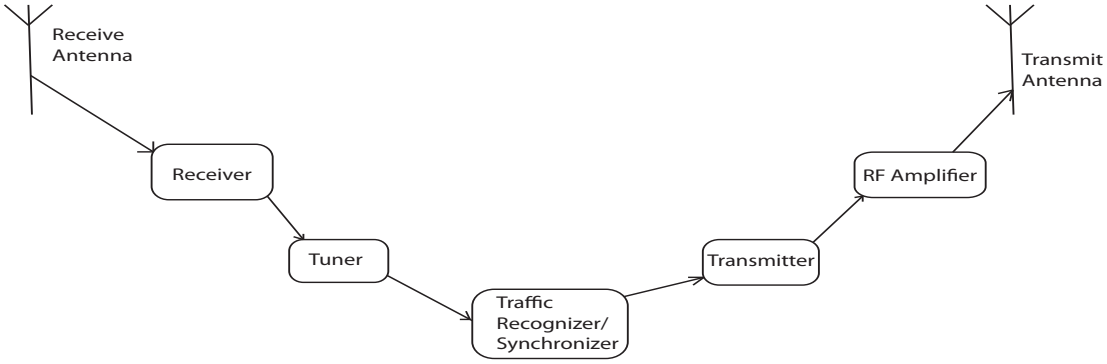

Figure 3: Simple Selective Jamming Architecture

5.4 Architecture

Recent work has shown that inexpensive software programmable radios such as the Ettus USRP are capable of implementing the P25 protocols and acting as part of a P25 deployment. Their versatility and the availability of open-source P25 software makes them attractive as well for use as jamming attack platforms. [5]

We are currently implementing a proof-of- concept low duty cycle selective P25 jammer on the USRP platform. See Figure 3. Our architecture is based on a USRP equipped with receiver and transmitter daughterboards and a small outboard RF amplifier. Each received signal is analyzed in real time by a recognizer filter as it arrives; if the received headers indicate a targeted transmission (e.g., one that is encrypted, associated with a particular NAC or TGID, or of a particular message type), the transmitter is pulsed in sync with the NID fields of each received frame. For voice traffic, the jammer operates at a duty cycle of 3.7% compared with the target transmitter.

6 Conclusions

APCO P25 is a widely deployed protocol aimed at critical public safety, law enforcement, and national security applications. The user base for secure P25 is rapidly growing in the United States and other countries, especially among federal law enforcement and intelligence agencies that conduct surveillance and other covert activities against sophisticated adversaries.

As a wireless system, P25 is inherently vulnerable to passive traffic interception and active attack, and so it must rely entirely on cryptographic techniques for its optional security features. And yet we found the protocols and its implementations suffer from serious weaknesses that leak sensitive data, invite inadvertent clear transmission in “secure” mode, and permit active and passive tracking and traffic analysis.

The protocol is particularly vulnerable to denial of service. Perhaps uniquely among modern digital radio systems, P25 systems can be effectively jammed with only a fraction of the aggregate signal power used by the legitimate user, by attackers with low cost equipment and without access to secrets such as keys or user-specific codes. Jamming attacks can also be used to aid in the exploitation of other weaknesses, such as selectively disabling security features to force users into the clear.

We are currently conducting a range of experiments with jamming and other attacks.

Acknowledgements

Partial support for this work was provided by a grant from the National Science Foundation, CNS-0905434.

References

[1] Leemon C. Baird III, William L. Bahn, Michael, and D. Collins. Jam-resistant communication without shared secrets through the use of concurrent codes, 2007.

[2] Stephen Bartlett. Does the digital radio standard come up short?, "April" 2001. http://urgentcomm.com/mag/radio_digital_radio_standard/.

[3] C. Cook and H. Marsh. An introduction to spread spectrum. *Communications Magazine, IEEE*, 21(2):8 – 16, March 1983.

[4] Daniels Electronics. P25 Training Guide, 2009. http://www.danelec.com/library/english/p25_training_guide.asp.

[5] Stephen Glass, Marius Portmann, and Muthukkumarasamy Vallipuram. A software-defined radio receiver for apco project 25 signals. In *International Workshop on Advanced Topics in Mobile Computing for Emergency Management: Communication and Computing Platforms*, pages 67–72, Leipzig, Germany, May 2009. ACM.

[6] Wang Hang, Wang Zanji, and Guo Jingbo. Performance of dsss against repeater jamming. In *Electronics, Circuits and Systems, 2006. ICECS '06. 13th IEEE International Conference on*, pages 858 –861, dec. 2006.

[7] Nathaniel Husted and Steven Myers. Mobile location tracking in metro areas: malnets and others. In *Proceedings of the 17th ACM conference on Computer and communications security*, CCS '10, pages 85–96, New York, NY, USA, 2010. ACM.

[8] H. T. Kung and D. Vlah. Efficient location tracking using sensor networks. In *Wireless Communications and Networking, 2003. WCNC 2003. 2003 IEEE*, volume 3, pages 1954–1961 vol.3. IEEE, 2003.

[9] J McCune. Dsss vs fhss narrowband interference performance issues, 2000.

[10] Motorola. Motorola P25 Compliance. http://esp.ongov.net/OCICS/documents/Motorola_P25_Compliant_Features.pdf.

[11] Motorola. Motorola-USA Digital Portable Radios. http://www.motorola.com/Business/US-EN/Business+Product+and+Services/Two-Way+Radios+-+Public+Safety/P25+Portable+Radios/XTS5000_US-EN.

[12] Ryan Naraine and ZDNET.com. Metasploit's hd moore releases 'war dialing' tools, March 2009. http://www.zdnet.com/blog/security/metasploits-hd-moore-releases-war-dialing-tools/2808.

[13] Roger L. Peterson, David E. Borth, and Roger E. Ziemer. *An Introduction to Spread-Spectrum Communications*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1st edition, 1995.

[14] Telecommunications Industry Association. APCO Project 25 - Over-the-Air-Rekeying(OTAR) Protocol. Technical Report TIA-102.AACA. http://www.tiaonline.org/standards.

[15] Telecommunications Industry Association. Project 25-DataOverview-NewTechStandards. Technical Report TIA-102.BAEA-A. http://www.tiaonline.org/standards.

[16] Telecommunications Industry Association. Project 25-Vocoder Description Standard. Technical Report TIA-102.BABA. http://www.tiaonline.org/standards.

[17] Trad, Baird, and Mitola. Software radio architecture evolution- foundations, technology tradeoffs, and architecture implications, 2000.